

ABSTRACT

A scaled-down representation of input to a compute-intensive application is created. A computing requirement based on the scaled-down representation is calculated. A turn-around time and an actual cost to a customer to run the compute-intensive application with the input, on one or more processors, based on the calculated computing requirement, is calculated and then sent to the customer.